

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION FOR U.S. LETTERS PATENT

Title:

GRANULAR ACCESS CONTROL OF INTER-PROCESS COMMUNICATIONS IN A
COMPARTMENT MODE WORKSTATION LABELED ENVIRONMENT

Inventors:

Scott Alan Leerssen
4282 Roswell Road, G12
Atlanta, GA 50342-3726

Citizenship: U.S.A.

Paul Anthony Cooke
356 Kinross Drive
Walnut Creek, CA 94598

Citizenship: U.S.A.

Suresh Ganesh Pai
3578H Meadowglen Village Lane
Doraville, GA 30340

Citizenship: India

Janak Ratilal Desai
66 Saint Claire Lane
Atlanta, GA 30324

Citizenship: U.S.A.

**GRANULAR ACCESS CONTROL OF INTER-PROCESS COMMUNICATIONS IN A
COMPARTMENT MODE WORKSTATION LABELED ENVIRONMENT**

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This application is a continuation of Application Serial No. 09/418,285, entitled “GRANULAR ACCESS CONTROL OF INTER-PROCESS COMMUNICATIONS IN A COMPARTMENT MODE WORKSTATION LABELED ENVIRONMENT”, filed October 14, 1999, the disclosure of which is hereby incorporated herein.

FIELD OF THE INVENTION

[0002] The present invention relates to a method for maintaining a secure operating system runtime environment and, more particularly, to a secure operating system having defined arbitrary relationships between subjects and objects of different sensitivity labels.

DESCRIPTION OF RELATED ART

[0003] A Compartment Mode Workstation (CMW) is a secured operating system that meets specific requirements set forth by the U.S. Government. Mandatory Access Control (MAC) and privileges are two important elements that go into creating a CMW.

[0004] MAC is a system-enforced access control mechanism that uses clearances and sensitivity labels to enforce security policy. MAC associates a security level (which may be either a defined clearance or sensitivity label) with every entity (passive and active) on the system. MAC permits read access to other entities at the same or lower security level, but prevents entities from writing to a lower level. In other words, an entity can read and write a subject or object at the same level but can only read an entity at a lower level. An active entity is defined as a subject, for example a process which causes information to flow amongst objects or changes the system state. A passive entity is defined as an object, and is typically “acted” upon. MAC is always enforced and cannot be overridden without privilege.

[0005] A CMW keeps track of security levels with clearances. A clearance represents the degree of security with which an entity is entrusted and has two components:

Classification - a hierarchical level of security. When applied to a process, the classification represents a measure of trust; applied to data, it is the degree of protection required

by the data.

Compartment - a subdivision of a classification that represents a grouping. Access is generally granted on a need-to-know basis.

[0006] A sensitivity label is an ASCII string representation of the classification/compartment combination of an objects security level. All subjects and objects in a CMW system have sensitivity labels. Since a subject is an active entity, it usually causes information to flow amongst objects or changes the system state. An object is generally a passive entity that contains or receives data, such as files, devices, printers, network interfaces, etc.

[0007] A CMW mediates all attempted security-related transactions by comparing the subject's sensitivity label with the object's sensitivity label. It permits or disallows the transaction dependent upon which label is dominant. An entity's sensitivity label dominates another if both of the following conditions are met:

- 1) The classification of the first entity's sensitivity label is equal to or outranks the object's classification.
- 2) All compartments in the first entity's sensitivity label are included in the second's sensitivity label.

[0008] Two labels are equal if they have the same classification and the same set of compartments. If they are equal, they dominate each other - thus access is permitted. If one label has a higher classification or includes all of the second label's compartments or both, the first label strictly dominates the second. Two labels are considered incomparable if neither label dominates the other.

[0009] MAC permits a subject to write up information and to read down information. In a read operation, the subject's sensitivity label must dominate the object's sensitivity label. Thus ensures that the subject's level of trust meets the requirements for access to the object and the subject's sensitivity label includes all compartment groupings that are allowed access to the object. In a write operation, the resulting object's sensitivity label must dominate the subject's sensitivity label. This prevents the subject from lowering the object's sensitivity label.

[0010] There are times when a security policy needs to be overridden. On conventional systems there is generally a single super-user account which has the ability to override all security policy. On a CMW system, the all powerful super-user authorization checks are replaced by individual checks for discrete privileges. Such a mechanism allows an administrator the ability to give a subject only those privileges that are required to perform its required task, and not give it complete override capability over the system's security policy.

[0011] U.S. Patent No. 5,903,732 entitled "Trusted Gateway Agent For Web Server Programs", issued on May 11, 1999 and hereby incorporated by reference, discloses an invention incorporated into the Virtual Vault ("VV") product sold by Hewlett-Packard Company. This product relies primarily upon MAC and privilege to provide a secure runtime environment for arbitrary Internet aware applications. VV separates application components and assigns each component a unique sensitivity label. In general, an application is divided into three unique components:

1) A read only component with a low classification on the system. This is usually the application's configuration files, libraries and executables.

2) An Internet component. The sensitivity label of this component dominates the read only component and is incomparable to the Intranet component. This component usually consists of application subjects that handle requests from the Internet.

3) An Intranet component. The sensitivity label of this component dominates the read only component and is incomparable to the Internet component. This component usually consists of application subjects that access information from the Intranet.

[0012] The Internet and Intranet components have incomparable levels and cannot communicate with each other. A trusted mechanism (Trusted Gateway Agent - "TGA") provides for communication between subjects and objects at incomparable levels. For example, the TGA mediates access between a web server (Internet component) and Common Gateway Interface (CGI) applications (Intranet component).

[0013] While traditional privileges and dedicated proxies can be used, they each suffer from their own particular problems. A subject that uses one of the CMW privileges has complete access to all other subjects and objects on the system, including subjects that have a dominate sensitivity labels and all subjects which have incomparable sensitivity labels. This privilege cannot restrict communications between specific components and thus is less secure.

The use of dedicated proxies, while allowing finer grained access controls between components, is ill suited to the Internet environment where the number of deployed protocols is large and continues to grow. Developing a separate proxy for each protocol would be impractical.

[0014] It would be desirable to provide for granular access control of inter-process communications in a CMW environment. Provider access control overcomes the traditional privilege problem when a user with CMW privilege has complete access to all other subjects on the system while allowing for the deployment of single and multi-tier Internet applications without the explicit need for dedicated proxies.

[0015] It would be desirable and of considerable advantage to provide a secure operating system that differs from that employed in prior art MAC based CMW systems. Further, such a secure operating system could be advantageous when implemented by use of sensitivity labels comprised of security levels and compartments, especially if the new secure operating system restricts/prohibits the transfer of data between subjects and objects of differing security levels and the security levels are used to label all system subjects and objects including, but not limited to; network connections,-file system objects, processes, shared memory and message IPC.

[0016] It will be apparent from the foregoing that there is still a need to define arbitrary relationships between objects including but not limited to filesystem objects, shared memory/messages, TCP and tm so as to provide discrete access between subjects and objects having arbitrary, normally incomparable sensitivity labels.

BRIEF SUMMARY OF THE INVENTION

[0017] The present invention provides a method for maintaining a secure run-time operating system environment in which sensitivity labels are enforced such that the operating system restricts the transfer of data between subjects and objects of differing sensitivity labels where a sensitivity label must dominate or be considered incomparable to other sensitivity labels.

[0018] Arbitrary relationships between subjects and objects of differing sensitivity labels are defined wherein the secure operating system enforcement of sensitivity label dominance and the secure operating system prohibition of data transfer between incomparable sensitivity labels is extended to provide discrete access between arbitrary, normally

incomparable sensitivity labels. Other aspects and advantages of the present invention will be come apparent from the following detailed description, taken in conjunction with the accompanying drawings, illustrating by way of example the principles of the invention.

BRIEF DESCRIPTION OF THE DRAWINGS

[0019] Figure 1 is a block diagram illustrating the application of privileges to establish IPC channels.

[0020] Figure 2 is a block diagram illustrating the establishment of IPC channels between host machines having unique sensitivity labels.

[0021] Figure 3 is a block diagram of the data-partitioning format employed by the invention which employs a fixed set of classifications and a dynamic set of compartments.

[0022] Figure 4 is a block diagram illustrating the cross boundary IPC mechanism which mediates access between privileged processes at one sensitivity label with subjects and objects at another sensitivity label.

[0023] Figure 5 is a flow chart illustrating the method steps of the claimed invention.

DETAILED DESCRIPTION

[0024] It would be desirable to provide a flexible and configurable secure operating system run-time environment that allows MAC to be overridden for inter-process communications between subjects and objects having different sensitivity labels. In particular, when a prior art MAC based operating system encounters a request from an application process to establish an IPC channel, it will ensure that MAC will not be invalidated if the channel is established. Only if the subject making the request is privileged would the OS allow MAC to be overridden. Such application of privilege allows for the establishment of IPC channels as shown in Figure I for traditional privileged processes.

[0025] The present invention employs a MAC based privilege with additional security checks employed to augment the authority of such privilege. When the secure OS in accordance with the invention encounters a request from an application process to establish an IPC channel, it will still ensure that MAC will not be invalidated if the channel is established. In

particular, for MAC to be overridden and the IPC channel to be successfully instantiated the following three conditions must be met:

- (1) The subject making the IPC request must be privileged.
- (2) The subject or object that is the target of the request must be privileged.
- (3) There must be a configured relationship between the sensitivity label of the requesting subject and the target subject or object.

[0026] Such application of privilege allows for the establishment of IPC channels as shown in Figure 2. Through granular access control, this mechanism provides a secure run time environment for application processes that is more secure than traditional CMW systems. The invention provides that each network interface of the host machine is assigned a unique sensitivity label (VAULT OUTSIDE and VAULT INSIDE).

[0027] As illustrated in Figure 3, the invention provides a data-partitioning scheme having a fixed set of classifications and a dynamic set of compartments. This partitioning provides complete information separation between VV components, network interfaces, each application's content and every deployed application component. The set of classifications provides MAC separation between the VV files, the deployed application content files and the application processes. The SYSTEM classification provides a unique compartment to hold all operating system files and the read-only files of the VV layered components. The CONTENT classification holds all application deployed read-only data files. The VAULT classification is used to run all of the Internet applications and to hold any application deployed read-write data files. These extra classifications provide for another layer of defense against undesirable modification of files and directories, and a more secure infrastructure for allowing multiple application integrations and content manipulation on a single machine.

[0028] Figure 4 illustrates the preferred embodiment of the invention having a flexible and configurable cross boundary IPC mechanism which mediates access between privileged processes at one sensitivity label with subjects and objects at another sensitivity label. These subjects and objects can be other privileged processes or network interfaces. This architecture isolates application components within their own compartments and provides for controlling access between components by mediating access between the compartments. For example, the invention provides for a privileged process within the VAULT WEB SERVER compartment access to the VAULT OUTSIDE network interface (Figure 4). It also allows a

privileged process within the CGI compartment access to the VAULT INSIDE network interface. The invention does not allow communications between a privileged process in the VAULT WEB SERVER compartment and a privileged process in the CGI compartment, for this inter-compartment mapping has not been configured. All of these cross-compartment communication relationships are configurable. Thus, the invention provides for a unique combination of a multi-classification, multi-compartment VV that uses a novel security mechanism to provide a secure deployment framework for a wide variety of Internet application architectures.

[0029] Mapping defines the allowed communications channels between privileged and non-privileged processes that do not possess the same sensitivity label. The ability to separately map individual processes with or without communication privilege and their communication mapping attributes is provided within each compartment. Therefore, a compartment to compartment communications mapping can be configured to perform an effective privilege access check by mapping attributes “with privilege” and to bypass the access check by mapping attributes “without privilege”.

[0030] Mapping is employed for defining arbitrary relationships between subjects and objects having different sensitivity labels. Once provided, discrete access between these arbitrary, normally incomparable sensitivity labels is enabled. In the preferred embodiment of the invention, basic mapping employs the arrow “==>” to represent data flow (for file mapping operations only). In addition to data flow, a mapping can be restricted. A restricted mapping requires the process to have a privilege. If a “P” is attached to the “==>”, then the process using the mapping must hold a special privilege to use the mapping.. For example, the mapping:

system outside (rc) P==> system inside

dictates that a system outside labeled process may connect to a system inside process/object or open for reading a system inside object, but it must hold a privilege to do so. Take away the “P”, and any system outside process can use the mapping. For example, the mapping of an outside network interface to an outside server process setup as safeclient and safeserver (Where safeclient is synonymous with (Send(UDP), Connect(TCP) and safeserver is synonymous with (Accept(TCP), Receive(UDP) in communications between compartments):

System Outside(A) P==> Outside NIC
Outside NIC(C) P==> System Outside

[0031] Another example of this mapping would be a mapping of an outside network interface to an outside server process setup with safeserver only:

System Outside(A) P==> Outside NIC

Outside NIC(C) => System Outside

[0032] The following mapping examples (without privilege mappings) illustrates the format of a valid mappings section where the attributes are the operations that the subjects on the left: side of the ==> can do to the objects on the right side. The attributes are set up as a Policy file within the operating system and include:

a - accept (TCP)

c - connect (TCP)

r - read

w -write

s - send (UDP)

r - receive (UDP)

Additional mappings include:

System Outside(attributes) ==> Vault Web Server

Vault Web CGI (attributes) ==> Vault Inside _NIC

System Outside(attributes) ==> Vault App Front End

Vault App Front End(attributes) ==> Vault App Back End

Vault App Back End(attributes) ==> Vault Inside_NIC

Alternatively, numeric mappings may be employed:

101(attributes) ==> 102

104(attributes) ==> 100

101(attributes) ==> 106

106 (attributes) ==> 107

107(attributes) ==> 100

[0033] Whenever the MAC daemon within a secure operating system reads and activates the contents of the Policy file, this simple decision list is loaded into the kernel and held in an existing communications decision cache. Keeping this information in the kernel eliminates the need to consult the MAC policy daemon for communications decisions.

[0034] The invention also provides for IPC mechanisms (message queues, semaphores and shared memory) restriction to the local machine and that they are free from any network interface complications. Each of these mechanisms use a similar API to both create a new IPC mechanism and to utilize an existing IPC mechanism. If a process uses one of the IPC APIs to create a new communications channel (msgget(), semget() or shmget() and has the safeserver privilege in its effective set, then the newly created IPC channel is marked that it was created by a safe server application.. (Where safeclient is synonymous with (Send(UDP), Connect(TCP) and safeserver is synonymous with (Accept(TCP), Receive(UDP) in communications between compartments). The difference occurs when a client program attempts to connect to an existing IPC channel with a different sensitivity label.. In the context of the invention, the IPC channel will be established if one the following conditions are met:

The attribute (safeclient or safeserver) controls are mapped with privilege (P ==>) and the access check for effective privilege by the system is successful.

The attribute (safeclient or safeserver) controls are mapped without privilege (==>).

The attribute Accept or Receive control is mapped and the appropriate network privilege access check is successful.

Communications within the invention can be bi-directional if one of the following conditions are met:

The attribute (safeclient or safeserver) controls are mapped in both directions between compartments with privilege (P ==>) and the access check for effective privilege by the system is successful.

The attribute (safeclient or safeserver) controls are mapped in both directions between compartments without privilege (==>).

The attribute Accept or Receive control is mapped in both directions and the appropriate network privilege access check is successful.

[0035] The invention can be configured in a compatibility mode by performing a bi-directional mapping of safeclient and safeserver controls between the OUTSIDE and INSIDE compartments with privilege. For illustrative purposes, the “vault app front end(C) P ==> vault app back end” and the reciprocal “vault app back end (A)P ==> vault app front end” mapping defined in the above example are employed. A process with the safeclient privilege and a

sensitivity label of Vault Front End can connect to a System V IPC channel created by a process with the safeserver privilege and a sensitivity label of Vault Back End.

[0036] The TCP(accept and connect) and UDP(send and receive) IPC mechanisms operate in the same way when the client and server reside on the same machine, but are modified if a network interface is involved. The architecture employed by the invention defines that network interfaces are implicitly unprivileged which means that they do not hold the safeserver and safeclient privileges as a default. This implies that network interfaces cannot connect or send to a processes unless explicitly given safeclient privilege or granted privilege by the new mappings. The implied privileges require some special consideration of data that arrives or attempts to leave the machine.

[0037] Where the network interface connected to the Internet has been labeled Vault Outside_NIC and the network interface connected to a corporate Intranet has been labeled Vault Inside_NIC, and TCP and a server process, which has a sensitivity label of Vault Web Server, the process creates a socket, binds it to a network port, and begins to listen for incoming connections. At some point, the process will have to call accept() to accept an incoming connection. If the incoming connection request is from a network interface, there must exist a mapping between the label of the network interface and the label of the process, Based on the supplied “vault web server(A) ==> vault outside_nic” mapping, the process under scrutiny could accept connections from the Internet but not from the Intranet, since no mapping has been defined for the later case. This is still a one way street as the mapping specifies a specific direction. Even if the server process held the safeclient privilege, there is no mapping from Vault Web Server to Vault Outside_NIC.

[0038] As another example, a process that has a sensitivity label of Vault Web CGI and holds the safeclient privilege attempts to make a connection to a remote machine on the Intranet. Since the connection would utilize the network interface with the Vault Inside_NIC sensitivity label, the “vault web CGI(C) P ==> vault inside_nic”, and a reciprocal “vault inside_nic(A) P==>vault web CGI” mapping, must exist for this connection request to continue.. Since the mapping is defined, connect filtering would allow the connection request to continue.

[0039] While the invention has been described and illustrated with reference to specific embodiments employing a UNIX based CMW (Compartment Mode Workstation), those

skilled in the art will recognize that modification and variations may be made such that the invention is equally applicable to other secure Web platforms and most compatible hardware that provide for defining arbitrary relationships between subjects and objects of differing sensitivity levels.